

09450001-120000

5           Generally, the present invention relates to memory system architecture and,  
in particular, the present invention relates to cache design.

The speed at which computer processors can execute instructions continues to outpace the ability of computer memory systems to supply instructions and data to the processors. Consequently, many high-performance computing systems provide a high-speed buffer storage unit, commonly called a cache or cache memory, between the working store or memory of the central processing unit (“CPU”) and the main memory.

A level 1 cache (“L1”) generally refers to a memory bank built closest to the  
30 central processing unit (“CPU”) chip, typically on the same chip die. A level 2



















5

10

15

20

25

30

The apparatus and method of example embodiments of the invention  
30 prevents premature eviction of cache lines allocated to a cache by a run-ahead









other data that is no longer needed by the processor and may be overwritten. A potential victim may be protected or unprotected, as indicated by its associated protection bit.

In one embodiment, the method 300 further comprises allocating a cache line into the cache to replace the potential victim 308, and setting a protection bit associated with the allocated cache line 310. Lines allocated during run-ahead execution are protected from eviction. This prevents the run-ahead prefetcher from evicting earlier prefetched lines that will be useful, once normal execution is resumed. This also prevents the run-ahead prefetcher from running too far ahead.

10 In one embodiment, the run-ahead prefetcher is capable of executing about 10,000 instructions while waiting for a memory reference.

In one embodiment, the method 300 further comprises switching to normal execution 312, referencing the allocated cache line 314, and clearing the protection bit associated with the allocated cache line 316. At some point during run-ahead execution the processor may switch to normal execution. For example, once the data is retrieved for the cache miss that initiated run-ahead execution, the processor switches to normal execution. Then, when a cache line is referenced by normal execution, its protection bit is cleared so that it is unprotected and free to be used in future allocations by the run-ahead prefetcher. Thus, clearing protection bits makes

15 room in the cache for more run-ahead prefetching.

Figures 4A-4C show flow charts of example embodiments of a method of accessing a cache. One aspect of the present invention is a method of accessing a cache 400, shown in Figure 4A, comprising determining whether a mode is run-ahead execution or normal execution 402, and replacing a first cache line 408 upon a cache miss 404 during run-ahead execution only if a protection bit associated with the first cache line is clear 406. Run-ahead allocation protection prevents cache lines prefetched earlier by the run-ahead prefetcher from being evicted as well as preventing cache lines currently in use by normal execution from being evicted. These cache lines are protected by protection bits.

25





are referenced, they are unprotected and free to be used again. In this way, the software prefetching thread produces cache entries that are consumed during normal execution.

In one embodiment, the method 500 further comprises clearing all protection bits when the software prefetching thread finishes executing 508. In one embodiment, the method 500 further comprises spawning the software prefetching thread for a predetermined section of code in the program 510. In one embodiment, the method 500 further comprises providing code for a software prefetching thread from an optimizing compiler 512.

10           An example embodiment of the present invention is illustrated by  
pseudocode shown in Table 1. This method, which is invoked for each cache  
access, is a technique for preventing thrashing in a cache augmented with a run-  
ahead prefetcher. Experiments have shown that this technique is successful at  
preventing the cache from thrashing even in the presence of very aggressive  
15   prefetchers. As memory latencies grow to several thousand instructions,  
independently sequenced prefetchers will become more common. As run-ahead  
prefetchers become more common, this technique to prevent the independently  
sequenced prefetcher from getting too far ahead of the program's thread of  
execution will be even more useful.

Table 1

```

5  struct cache_line_struct{
    unsigned long tag; /* line tag */
    char valid; /* valid bit */
    char dirty; /* dirty bit */
    char protected; /* protection bit */
10  char *data; /* line data */
    } line;

    struct cache_set_struct{
        line lines[NUM_ASSOC]; /* lines[0] = LRU, lines[NUM_ASSOC-1] = MRU */
15  } cache_set;

    struct cache_struct{
        cache_set sets[NUM_CACHE_SETS];
    } cache;
20  cache c;

    char* /* return line data */
    do_cache_access(unsigned long addr, /* address of access */
25  int TOA, /* type of access RD/WR */
        int run_ahead, /* 1 = run ahead mode, 0 = normal mode */
        char* data /* for writes */
        ) {
        unsigned long tag = GET_TAG(addr);
30  unsigned int set = GET_SET(addr);
        unsigned line *l = find_line_in_set(tag, c.sets[set]);
        unsigned line *repl;

        if (!run_ahead) {
35  if (l) { /* if a hit */
            l->protected = 0;
            update_LRU(c.sets[set], l); /* place l at the head of LRU list */
            if (TOA == RD) /* read */
                return l->data;
40  else { /* write */
            l->dirty = 1;
            return l->data = data;
        }
    } else { /* miss */
45  repl = &(c.sets[set].lines[0]); /* replace LRU block */
        process_miss(addr, TOA, run_ahead, data.repl);
    }
}

```







000221" 02051460

cache controller 716 is associated with the plurality of caches 708 (shown in Figures 7 and 8). In this embodiment, the control logic 714 resides in the at least one cache controller 716. However, all or part of the control logic 714 may reside elsewhere.

In one embodiment, the multiprocessor computer system 700 further comprises a plurality of tag arrays 718, as shown in Figure 10. A tag array 718 is associated with each cache 708 (shown in Figures 7 and 8). In this embodiment, the protection bits 712 reside in each tag array 718 and are associated with cache lines 710. A tag is the remainder of an address generated by the processor after the set bits have been removed. Set bits are the address used to find a line within a cache. The cache management logic may compare the tag bits of the address with the tag bits of the cache directory which are stored at the same set address.

One aspect of the present invention is a computer system comprising a main memory, a processor, a bus, a cache, and a protection bit. The computer system may be any system including, but not limited to, the systems shown in Figures 1, 6A-6C, 7, or 8. The bus connects the main memory and the processor. The cache is associated with the processor and has a plurality of cache lines. The protection bit is associated with each of the cache lines in each of the plurality of caches. Each protection bit protects a cache line from premature eviction during speculative execution. In one embodiment, the cache is a level one (L1) cache and in another embodiment, the cache is a level two (L2) cache. In one embodiment, the L1 cache is on the same chip die as the processor.

It is to be understood that the above description it is intended to be illustrative, and not restrictive. Many other embodiments will be apparent to those of skill in the art upon reviewing the above description. The scope of the invention should, therefore, be determined with reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.